

CLAIMS

1. A method of processing data in a computer system comprising at least one host and at least one storage system, the method comprising acts of:

(A) receiving a request, from the host, to delete a unit of data stored on the storage system;

(B) determining whether a previously-defined retention period for the unit of data has expired; and

(C) when it is determined in the act (B) that the retention period for the unit of data has not expired, denying the request to delete the unit of data.

2. The method of claim 1, further comprising an act of:

(D) when it is determined in the act (B) that the retention period for the unit of data has expired, deleting the unit of data.

3. The method of claim 1, wherein the acts (A), (B) and (C) are performed by the storage system.

4. The method of claim 1, further comprising an act (D) of, prior to performing the acts (A), (B) and (C), receiving information specifying the retention period for the unit of data.

5. The method of claim 1, further comprising acts of, prior to performing the acts (A), (B) and (C):

(D) receiving the unit of data at the storage system; and

(E) writing the unit of data to the storage system.

6. The method of claim 5, further comprising acts of, prior to performing the acts (A), (B) and (C):

(F) receiving information specifying the retention period for the unit of data along with the unit of data; and

(G) writing the information specifying the retention period to the storage system.

7. The method of claim 1, wherein the storage system stores at least two different types of units of data including a blob of data and a content descriptor file (CDF) that corresponds to the blob, the blob and the CDF belonging to blob/CDF set, the CDF including metadata relating to the corresponding blob, and wherein the unit of data is the CDF so that:

the act (A) comprises an act of receiving a request to delete the CDF;

the act (B) comprises an act of determining whether a previously-defined retention period for the CDF has expired; and

the act (C) comprises an act of, when it is determined in the act (B) that the retention period for the CDF has not expired, denying the request to delete the CDF.

8. The method of claim 7, further comprising an act (D) of denying any request to delete the blob unless and until the retention period for the corresponding CDF has expired.

9. The method of claim 8, wherein the act (D) comprises an act of denying any request to delete the blob unless and until the corresponding CDF has been deleted.

10. The method of claim 7, wherein the CDF is a first CDF, wherein the blob/CDF set comprises a second CDF corresponding to the blob, wherein each of the first and second CDFs has a retention period, and wherein the method further comprises an act (D) of denying any request to delete the blob unless and until a longer of the retention periods for the first and second CDFs has expired.

11. The method of claim 1, wherein the storage system stores at least two different types of units of data including a blob of data and a content descriptor file (CDF) that corresponds to the blob and forms a blob/CDF set with the blob, the CDF including metadata relating to the corresponding blob, and wherein the unit of data is the blob so that:

the act (A) comprises an act of receiving a request to delete the blob;
the act (B) comprises an act of determining whether a previously-defined retention period for the blob has expired; and
the act (C) comprises an act of, when it is determined in the act (B) that the retention period for the blob has not expired, denying the request to delete the blob.

12. The method of claim 7, further comprising an act (D) of executing a garbage collection utility that deletes the blob, but not until after the retention period for the corresponding CDF has expired.

13. The method of claim 1, wherein the unit of data includes timestamp information identifying when the unit of data was stored to the storage system, wherein the act (B) comprises examining the timestamp information and wherein the method further comprises an act of sending to the at least one host time information from a source trusted by the storage system so that the at least one host can use the time information in creating the timestamp.

14. The method of claim 12, further comprising an act (E) of denying any request to delete the blob directly.

15. The method of claim 12, wherein the act (D) comprises an act of executing a garbage collection utility that deletes the blob in response to the corresponding CDF having been deleted so that the blob does not have a corresponding CDF.

16. The method of claim 7, wherein the CDF is a first CDF, wherein the blob/CDF set comprises a second CDF corresponding to the blob, wherein each of the first and second CDFs has a retention period, and wherein the method further comprises an act (D) of executing a garbage collection utility that deletes the blob, but not until after the expiration of the longer of the retention periods for the first and second CDFs.

17. The method of claim 16, wherein the act (D) comprises an act of executing a garbage collection utility that deletes the blob in response to both of the first and second CDFs having been deleted so that the blob does not have a corresponding CDF.

18. The method of claim 16, further comprising an act (E) of denying any request to delete the blob directly.

19. The method of claim 9, wherein the CDF is a first CDF, and wherein the method further comprises an act of extending a retention period for the blob by writing a second CDF to the storage system, the second CDF corresponding to the blob and having a longer retention period than the first CDF.

20. At least one computer readable medium encoded with instructions that, when executed on a computer system, perform a method, the computer system comprising at least one host and at least one storage system, the method comprising acts of:

(A) receiving a request, from the host, to delete a unit of data stored on the storage system;

(B) determining whether a previously-defined retention period for the unit of data has expired; and

(C) when it is determined in the act (B) that the retention period for the unit of data has not expired, denying the request to delete the unit of data.

21. The at least one computer readable medium of claim 20, wherein the method further comprises an act of:

(D) when it is determined in the act (B) that the retention period for the unit of data has expired, deleting the unit of data.

22. The at least one computer readable medium of claim 20, wherein the acts (A), (B) and (C) are performed by the storage system.

23. The at least one computer readable medium of claim 20, wherein the method further comprises an act (D) of, prior to performing the acts (A), (B) and (C), receiving information specifying the retention period for the unit of data.

24. The at least one computer readable medium of claim 20, wherein the method further comprises acts of, prior to performing the acts (A), (B) and (C):

- (D) receiving the unit of data at the storage system; and
- (E) writing the unit of data to the storage system.

25. The at least one computer readable medium of claim 24, wherein the method further comprises acts of, prior to performing the acts (A), (B) and (C):

- (F) receiving information specifying the retention period for the unit of data along with the unit of data; and
- (G) writing the information specifying the retention period to the storage system.

26. The at least one computer readable medium of claim 20, wherein the storage system stores at least two different types of units of data including a blob of data and a content descriptor file (CDF) that corresponds to the blob and forms a blob/CDF set with the blob, the CDF including metadata relating to the corresponding blob, and wherein the unit of data is the CDF so that:

- the act (A) comprises an act of receiving a request to delete the CDF;
- the act (B) comprises an act of determining whether a previously-defined retention period for the CDF has expired; and
- the act (C) comprises an act of, when it is determined in the act (B) that the retention period for the CDF has not expired, denying the request to delete the CDF.

27. The at least one computer readable medium of claim 26, wherein the method further comprises an act (D) of denying any request to delete the blob unless and until the retention period for the corresponding CDF has expired.

28. The at least one computer readable medium of claim 27, wherein the act (D) comprises an act of denying any request to delete the blob unless and until the corresponding CDF has been deleted.

29. The at least one computer readable medium of claim 26, wherein the CDF is a first CDF, wherein the blob/CDF set comprises a second CDF corresponding to the blob, wherein each of the first and second CDFs has a retention period, and wherein the method further comprises an act (D) of denying any request to delete the blob unless and until a last to expire of the retention periods for the first and second CDFs has expired.

30. The at least one computer readable medium of claim 20, wherein the storage system stores at least two different types of units of data including a blob of data and a content descriptor file (CDF) that corresponds to the blob and forms a blob/CDF set with the blob, the CDF including metadata relating to the corresponding blob, and wherein the unit of data is the blob so that:

the act (A) comprises an act of receiving a request to delete the blob;

the act (B) comprises an act of determining whether a previously-defined retention period for the blob has expired; and

the act (C) comprises an act of, when it is determined in the act (B) that the retention period for the blob has not expired, denying the request to delete the blob.

31. The at least one computer readable medium of claim 26, wherein the method further comprises an act (D) of executing a garbage collection utility that deletes the blob, but not until after the retention period for the corresponding CDF has expired.

32. The at least one computer readable medium of claim 31, wherein the method further comprises an act (E) of denying any request to delete the blob directly.

33. The at least one computer-readable medium of claim 20, wherein the unit of data includes timestamp information identifying when the unit of data was stored to the storage system, wherein the act (B) comprises examining the timestamp information and wherein the method further comprises an act of sending to the at least one host time

information from a trusted source so that the at least one can use the time information in creating the timestamp.

34. The at least one computer readable medium of claim 31, wherein the act (D) comprises an act of executing a garbage collection utility that deletes the blob in response to the corresponding CDF having been deleted so that the blob does not have a corresponding CDF.

35. The at least one computer readable medium of claim 26, wherein the CDF is a first CDF, wherein the blob/CDF set comprises a second CDF corresponding to the blob, wherein each of the first and second CDFs has a retention period, and wherein the method further comprises an act (D) of executing a garbage collection utility that deletes the blob, but not until after the expiration of the longer of the retention periods for the first and second CDFs.

36. The at least one computer readable medium of claim 35, wherein the act (D) comprises an act of executing a garbage collection utility that deletes the blob in response to both of the first and second CDFs having been deleted so that the blob does not have a corresponding CDF.

37. The at least one computer readable medium of claim 35, wherein the method further comprises an act (E) of denying any request to delete the blob directly.

38. The at least one computer readable medium of claim 28, wherein the CDF is a first CDF, and wherein the method further comprises an act of extending a retention period for the blob by writing a second CDF to the storage system, the second CDF corresponding to the blob and having a longer retention period than the first CDF.

39. The at least one computer readable medium of claim 28, wherein the CDF is a first CDF, and wherein the method further comprises an act of extending a retention period for the blob by writing a second CDF to the storage system, the second CDF corresponding to the blob and having a longer retention period than the first CDF.

40. A storage system for use in a computer system comprising at least one host and the storage system, the storage system comprising:
at least one storage device to store data received from the at least one host; and
at least one controller that;
receives a request, from the host, to delete a unit of data stored on the storage system;
determines whether a previously-defined retention period for the unit of data has expired; and
when it is determined that the retention period for the unit of data has not expired, denies the request to delete the unit of data.

41. The storage system of claim 40, wherein the controller, when it determines that the retention period for the unit of data has expired, deletes the unit of data.

42. The storage system of claim 40, wherein the at least one controller receives information specifying the retention period for the unit of data.

43. The storage system of claim 40, wherein the at least one controller receives the unit of data and writes the unit of data to the at least one storage device.

44. The storage system of claim 43, wherein the at least one controller receives information specifying the retention period for the unit of data along with the unit of data and writes the information specifying the retention period to the at least one storage device.

45. The storage system of claim 40, wherein the storage system stores at least two different types of units of data including a blob of data and a content descriptor file (CDF) that corresponds to the blob and forms a blob/CDF set with the blob, the CDF including metadata relating to the corresponding blob, and wherein the unit of data is the

CDF so that when the at least one controller determines that the retention period for the CDF has not expired, the at least one controller denies the request to delete the CDF.

46. The storage system of claim 45, wherein the at least one controller denies any request to delete the blob until the retention period for the corresponding CDF has expired.

47. The storage system of claim 46, wherein the at least one controller denies any request to delete the blob unless and until the corresponding CDF has been deleted.

48. The storage system of claim 40, wherein the storage system stores at least two different types of units of data, including a blob of data and a content descriptor file (CDF) that corresponds to the blob, the CDF including timestamp information identifying when the CDF was created and the storage system sends time information based upon a time of the storage system to the host for use in creating the timestamp information.

49. The storage system of claim 45, wherein the CDF is a first CDF, wherein the blob/CDF set comprises a second CDF corresponding to the blob, wherein each of the first and second CDFs has a retention period, and wherein the at least one controller denies any request to delete the blob unless and until a longer of the retention periods for the first and second CDFs has expired.

50. The storage system of claim 40, wherein the storage system stores at least two different types of units of data including a blob of data and a content descriptor file (CDF) that corresponds to the blob and forms a blob/CDF set with the blob, the CDF including metadata relating to the corresponding blob, and wherein the unit of data is the blob so that when the at least one controller determines that the retention period for the blob has not expired, the at least one controller denies the request to delete the blob.

51. The storage system of claim 45, wherein the at least one controller executes a garbage collection utility that deletes the blob, but not until after the retention period for the corresponding CDF has expired.

52. The storage system of claim 51, wherein the at least one controller denies any request to delete the blob directly.

53. The storage system of claim 51, wherein the at least one controller executes a garbage collection utility that deletes the blob in response to the corresponding CDF having been deleted so that the blob does not have a corresponding CDF.

54. The storage system of claim 45 wherein the CDF is a first CDF, wherein the blob/CDF set comprises a second CDF corresponding to the blob, wherein each of the first and second CDFs has a retention period, and wherein the at least one controller executes a garbage collection utility that deletes the blob, but not until after the expiration of the longer of the retention periods for the first and second CDFs.

55. The storage system of claim 54, wherein the at least one controller executes a garbage collection utility that deletes the blob in response to both of the first and second CDFs having been deleted so that the blob does not have a corresponding CDF.

56. The storage system of claim 54, wherein the at least one controller denies any request to delete the blob directly.

57. A storage system for use in a computer system comprising at least one host and the storage system, the storage system comprising:
at least one storage device to store data received from the at least one host; and
at least one controller that;

receives at least one request, from the host, to store a unit of data in the storage system until at least the expiration of a retention period specified in the at least one request;

stores the unit of data in the at least one storage device; and

stores the information specifying the retention period in the at least one storage device.

58. A method of processing data in a computer system comprising at least one host and at least one storage system, the method comprising an act of:

(A) transmitting at least one request, from the host to the storage system, requesting that the storage system store a data unit until at least the expiration of a retention period specified in the at least one request.

59. The method of claim 58, wherein the storage system stores at least two different types of data units including a blob of data and a content descriptor file (CDF) that corresponds to the blob and forms a blob/CDF set with the blob, the CDF including metadata relating to the corresponding blob, the storage system denying any request to delete the blob until the retention period for the corresponding CDF has expired, and wherein the data unit is a blob so that the act (A) comprises an act of transmitting at least one request requesting that the storage system store the blob until at least the expiration of a retention period specified in the at least one request.

60. A method of processing data in a computer system comprising at least one host and at least one content addressable storage system that stores at least two different types of data units including a blob of data and a content descriptor file (CDF) that corresponds to the blob and forms a blob/CDF set with the blob, the CDF including metadata relating to the corresponding blob, each one of the blobs and CDFs being accessible to the at least one host computer using an identifier that is generated based on a content of the one of the blobs and CDFs, wherein the storage system executes a garbage collection utility that deletes blobs but only after the corresponding CDF has been deleted so that the blob does not have a corresponding CDF, the method comprising an act of:

providing the host with at least two options for deleting at least one of a blob and CDF in a blob/CDF set, the at least two options differing in terms of how much processing of the deletion is performed upfront in response to a deletion request and how much processing of the deletion is left for the garbage collection utility.

61. The method of claim 1, wherein the retention period is based, at least in part, on a creation time when the unit of data is written to the storage system, and wherein the method further comprises an act of, when the retention period is specified to the storage system, checking the creation time against time information accessible to the storage system.

62. The at least one computer-readable medium of claim 20, wherein the retention period is based, at least in part, on a creation time when the unit of data is written to the storage system, and wherein the method further comprises an act of, when the retention period is specified to the storage system, checking the creation time against time information accessible to the storage system.

63. The storage system of claim 40, wherein the retention period is based, at least in part, on a creation time when the unit of data is written to the storage system, and wherein storage system checks the creation time against time information accessible to the storage system when the retention period is specified to the storage system.

64. The method of claim 58, wherein the act (A) comprises transmitting at least one request that specifies the retention period based, at least in part, on a creation time when the request is transmitted, and further comprising acts of:

(B) receiving from a source specified by the at least one storage system, time information; and

(C) using the time information in establishing the creation time.